Available on CMS information server

CMS NOTE -2024/006

The Compact Muon Solenoid Experiment

CMS Note

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

17 July 2024 (v2, 18 July 2024)

Low transverse-momentum hadronic tau lepton reconstruction performance in the Run 3 Scouting dataset

The CMS Collaboration

Abstract

A technique for the reconstruction of low transverse momentum (p_{T}) hadronically decaying tau leptons in the CMS Run 3 Scouting datastream is presented. The addition of hadronic taus to the set of physics objects in Scouting opens possibilities for searches in new regions of phase space. This note summarizes the performance and result of the effort to create the Scouting τ with a p_{T} threshold of 5 GeV. An observation of the $\mathrm{Z} \to \tau\tau$ peak is shown for the first time in data Scouting.

1 Introduction

The tau lepton (τ), with its unique properties and decay modes, plays a crucial role in the search for new physics beyond the Standard Model (BSM). Most searches using τ leptons utilize the hadronic decays of the τ (τ_h) to enhance sensitivity [1]. However most current searches that use τ_h focus on high-momentum (p_{T}) objects [2]. In this note we describe a method to reconstruct low-momentum τ_h using the CMS detector at the LHC [3], which will help extend the search regions to lower BSM masses and lower transverse momenta, opening up a previously inaccessible region of phase space. This can potentially reveal signals of new particles or interactions that might be elusive at higher mass scales or higher τ momenta.

Although the LHC provides collisions every 25 ns, the CMS detector and data acquisition systems do not have the bandwidth to record every event. Instead, events of interest are selected using a two-tier trigger system. The first level (L1), composed of custom hardware processors, uses information from the calorimeters and muon detectors to select events at a rate of around 100 kHz within a fixed latency of about 4 μs [4]. The second level, known as the High-Level Trigger (HLT), consists of a farm of processors running a version of the Particle Flow (PF) algorithm [5] optimized for fast processing and reduces the event rate to around 1 kHz before data storage [6]. If the partial reconstruction by the HLT indicates that the event has passed some specific requirements, the event is fully reconstructed and stored for offline analysis, e.g., if the scalar sum of the transverse momenta of jets in the event (H_{T}) is above some threshold.

To handle the high rate of standard model events produced via the strong interaction, referred to as quantum chromodynamics (QCD) multijet events, the standard trigger selections require very high p_{T} thresholds for the reconstructed τ leptons. This is an obstacle in capturing low momentum τ_h. The CMS data scouting stream can be used to capture these events, where the reconstruction performed by the HLT is saved for further analysis [6]. Data scouting allows CMS to collect data at a rate much higher than possible with standard triggers. The data size of scouting events is approximately 100 times smaller than that of standard events. Scouting data exploits the considerable computing capability of the CMS data acquisition and filtering system [6] to reconstruct events at a high rate. Scouting data does not contain the detailed event level information available to the standard τ_h reconstruction algorithms. The complete list of physics objects available in the Scouting dataset are given in Ref. [6]. We use these objects to reconstruct τ_h candidates. In Run 3 of the LHC, the Scouting data stream at CMS features an almost complete menu of physics objects collected with low p_{T} triggers. Electrons, muons, photons, jets, and missing transverse energy (MET) are all successfully reconstructed at the HLT with resolutions close to those achieved by the standard offline reconstruction [6]. However, one critical physics object has been conspicuously absent from this lineup: hadronic τ leptons.

To address this gap, we developed the Scouting τ_h object with the aim of expanding the phase space available for analyses involving τ leptons in their final state. The development process involved meticulous adjustments to the reconstruction algorithm, with a particular focus on reconstructing low p_{T} τ_h objects. A minimum visible p_{T} threshold of 5 GeV was achieved for the τ_h object, compared to more than 20 GeV for the τ_h objects used in the standard offline reconstruction. Achieving this required several refinements in the reconstruction techniques to ensure accurate and efficient performance.

Extensive validation checks have been performed to verify the reliability and effectiveness of the Scouting τ_h reconstruction. Among these checks, we observe the $Z \to \tau_\mu \tau_h$ process with the expected cross section, for the first time using Scouting data. These observations confirm that the developed algorithm can successfully identify and reconstruct τ_h candidates, paving the way for enhanced analyses of processes involving low-p_T τ with the CMS experiment.

With the integration of τ_h reconstruction, the Scouting data stream now offers a comprehensive suite of physics objects collected with lower momentum thresholds than standard analysis data streams. This development not only broadens the scope of potential analyses, but also strengthens the overall scientific output of the CMS collaboration, enabling more precise and diverse investigations into the fundamental nature of particle physics.

2 Scouting τ_h

The τ lepton decays via either leptonic or hadronic channels. Leptonic decays require the reconstruction of electrons or muons. Hadronic τ decay produces combinations of charged and neutral hadrons [7]. The possible final states of τ_h are summarized in Table 1. The Scouting dataset includes Particle Flow (PF) [5] candidates of the following types: e/γ, K_L^0, $\pi^{\pm}$ and $\mu^{\pm}$. These reconstructed particles can be assembled into composite objects that meet the composition and kinematic characteristics consistent with τ_h decay. Charged PF particles with pixel-only tracks are stored with basic track parameters and are linked to the vertex associated with them. It should be noted that PF e/γ candidates are not separated by e or γ and do not include track information. All PF candidates in Scouting are limited to a $p_T > 0.6$ GeV and $|\eta| < 3$ [6].

Table 1: Summary of τ decay category types and branching ratios.

Decay mode	Branching Ratio [%]	Tag name
$\tau^- \to e^- \bar{\nu}_e \nu_\tau$	(17.82 ± 0.04)	
$\tau^- \to \mu^- \bar{\nu}_\mu \nu_\tau$	(17.39 ± 0.04)	
$\tau^- \to h^- \nu_\tau$	(11.51 ± 0.05)	1-prong
$\tau^- \to h^- h^+ h^- \nu_\tau$	(9.80 ± 0.05)	3-prong
$\tau^- \to h^- \pi^0 \nu_\tau$	(25.93 ± 0.09)	1-prong, 1-strip
$\tau^- \to h^- \pi^0 \pi^0 \nu_\tau$	(9.31 ± 0.05)	1-prong, 2-strip
Other	8.24	

Hadronically decaying τ leptons are reconstructed in the standard data stream via the hadrons-plus-strips (HPS) algorithm [8]. The algorithm uses charged and neutral pion candidates to combinatorially assemble τ_h decays. This study modifies the HPS algorithm to facilitate the construction of extremely low p_T τ_h decays using objects available in the Run 3 scouting dataset. The entire process is described below, including portions of the HPS algorithm that are unchanged from the standard implementation, to form a complete view.

2.1 Scouting HPS

2.1.1 Neutral hadron reconstruction

The neutral products (π^0) of τ_h decay promptly into two photons — these are reconstructed via PF e/γ candidates. The π^0 hypotheses are also called "strips" because the photons may pair-produce and form a vertical strip along ϕ by virtue of the magnetic field in the CMS detector. The HPS algorithm used in standard analyses utilizes an algorithm called dynamic strip reconstruction to allow better background rejection [8]. Scouting HPS employs the same algorithm as standard HPS but with objects available in Scouting. To cluster strips, the $\eta \times \phi$ plane is populated with e/γ candidates from an event. This plane is restricted to $|\eta| < 2.4$ and $-\pi < \phi < \pi$. The highest p_T e/γ candidate is selected as the seed for a strip. A rectangular box or strip is drawn around the seed. The dimensions of this strip are set by the following functions:

$$\begin{aligned}\Delta\eta &= f(p_{\mathrm{T}}^{e/\gamma}) + f(p_{\mathrm{T}}^{\mathrm{strip}}) \\ \Delta\phi &= g(p_{\mathrm{T}}^{e/\gamma}) + g(p_{\mathrm{T}}^{\mathrm{strip}})\end{aligned} \tag{1}$$

with

$$\begin{aligned}f(p_{\mathrm{T}}) &= 0.2 \cdot (p_{\mathrm{T}}[\mathrm{GeV}])^{-0.66} \\ g(p_{\mathrm{T}}) &= 0.35 \cdot (p_{\mathrm{T}}[\mathrm{GeV}])^{-0.71}.\end{aligned} \tag{2}$$

The next highest p_{T} e/γ candidate within the strip is added to the seed and the p_{T} weighted center is recalculated. A new strip is drawn around the recalculated center and the process repeats until no e/γ candidates can be found in the rectangle. In Eq. 1, $p_{\mathrm{T}}^{\mathrm{strip}}$ is the combined transverse momentum of the previously merged e/γ candidates in the strip. The next strip is seeded by the highest p_{T} e/γ candidate in the plane not yet merged into a strip. The Lorentz vector of every strip is rescaled to imbue the π^0 mass. Only strips with $p_{\mathrm{T}} > 1\,\mathrm{GeV}$ are considered in the next stages of Scouting HPS. The π^0 hypotheses are then reclustered into jets using FASTJET [9, 10] with the anti-k_{T} algorithm [11] and distance parameters ($\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$) of 0.4; these jets are referred to as AK4 jets. Strip constituents are removed from the jets to avoid double-counting of particles in the reclustering process.

2.1.2 Hadrons plus strips algorithm

The charged hadronic components of τ_h decay are taken directly from the PF Scouting particle collection. Reconstruction begins with the particle contents of an AK4 jet. Every possible combination of $\pi^\pm$ and π^0 in each jet consistent with a valid τ_h decay is considered. Within each jet, there may be many τ_h candidates with shared constituents. The τ_h candidate with the highest p_{T} in the jet is selected for that jet. At the event level, the most isolated τ_h is selected with isolation given by:

$$\tau_{iso} = \sum p_{\mathrm{T}}^{\mathrm{charged}}(d_z < 0.1\ \mathrm{cm}) + \max\left(0, \sum p_{\mathrm{T}}^{\gamma} - \Delta\beta \sum p_{\mathrm{T}}^{\mathrm{charged}}(d_z > 0.1\ \mathrm{cm})\right). \tag{3}$$

A cone of size $\Delta R = 0.4$ is considered when calculating isolation. That is, all particles within this cone that are not the τ_h itself are used to calculate the isolation. Here, the multiplicative factor to reduce pileup contribution to the isolation, $\Delta\beta$, is taken as 0.2, reflecting the Run 2 value [8]. The track longitudinal impact parameter with respect to the primary vertex is defined here as d_z. The p_{T} contribution to the isolation from charged particles is defined as $p_{\mathrm{T}}^{\mathrm{charged}}$ and that from e/γ candidates as p_{T}^{γ}. This is the only point in Scouting HPS where τ_{iso} is used. There are notable changes between the standard HPS and Scouting HPS algorithms. One effort made in the standard HPS algorithm to reduce QCD jets faking τ_h is the introduction of the signal cone and the τ cone. The signal cone is the region of the jet where the τ_h decay products are expected to be found. The size of the signal cone is calculated as $R_{\mathrm{signal\ cone}} = \frac{3.0}{E_T^{vis}[\mathrm{GeV}]}$ where E_T^{vis} is the transverse energy of τ_h and will confine the collimated decay products of the τ_h if it has sufficient p_{T}. The τ cone is the maximum ΔR between any of the τ_h constituents and the τ_h axis. If the τ cone is larger than the signal cone, then the τ_h candidate is rejected. Note that standard HPS does not allow the signal cone to be larger than $\Delta \mathrm{R} = 0.1$. For low p_{T} decays, the cone of the τ_h is almost always larger than 0.1, as shown in Figure 1 left. Note that the τ cone for the single prong τ_h is expressly zero. Scouting HPS introduces a dynamic signal cone limit such that at $p_{\mathrm{T}} > 20\,\mathrm{GeV}$, the limit matches standard HPS. At low p_{T}, there is an exponential cut made between the signal cone and p_{T}. The dynamic limit is given by Eq. 4 — this functional form and parameters therein were chosen to envelope the majority population of τ_h in $\Upsilon \to \tau_\mu \tau_h$ simulation.

$$\text{signal cone} < 2.8^{-0.3\times(p_{\mathrm{T}}(\tau_h)[\mathrm{GeV}]-1.4)+0.1}. \tag{4}$$

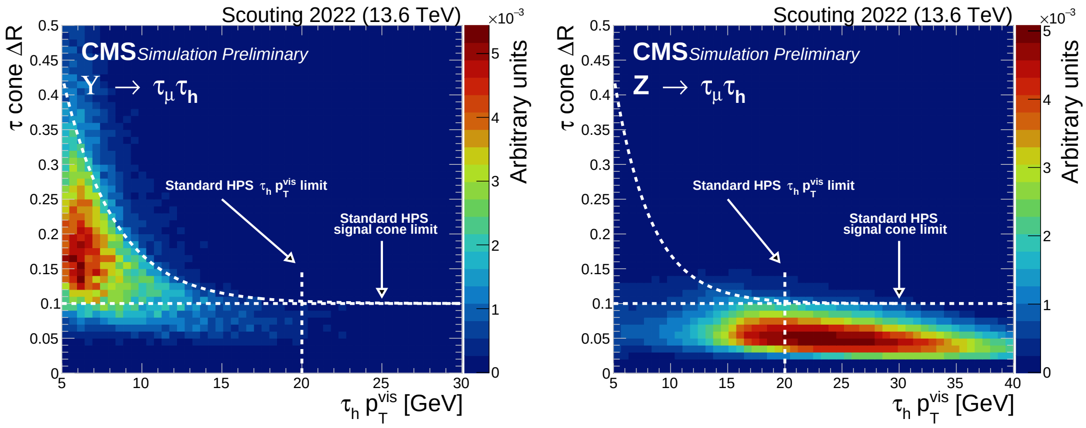

Figure 1: Generator-level τ cone ΔR (the maximum ΔR between the τ_h candidate axis and its constituents as determined from τ_h daughters at generator-level) plotted against τ_h p_{T} for $\Upsilon \to \tau_\mu \tau_h$ (left) and $Z \to \tau_\mu \tau_h$ (right) accepted by Scouting HPS without the cone requirement. Standard HPS limits the maximum signal cone to $\Delta R = 0.1$ (horizontal dotted line). For lower p_{T} τ_h candidates, it is necessary to open the allowed signal cone in order to increase sensitivity (curved dotted line). Note that the τ cone for 1-prong τ_h is expressly zero, and they are excluded from these plots.

There are numerous requirements summarized in Table 2 for each decay category that a candidate τ_h must pass to be considered in the jet-level pool.

Table 2: Summary of τ_h reconstruction requirements by decay category. Here $m_{vis}(\tau_h)$ is the visible invariant mass of the τ_h decay products.

Variable/Decay mode	1-prong, 1-strip	1-prong, 2-strip	1-prong	3-prong
τ_h visible mass [GeV]	$0.3 < m_{vis}(\tau_h) < 1.3$	$0.4 < m_{vis}(\tau_h) < 1.2$	—	$0.8 < m_{vis}(\tau_h) < 1.5$
Lead track p_{T} [GeV]	> 1.0	> 1.0	> 5.0	> 2.5
$\Delta R(\tau_h, \text{jet})$	< 0.1	< 0.1	< 0.1	< 0.1
Strip mass [GeV]	—	$0.05 < m < 0.2$	—	—

A simulation-level reconstruction efficiency study is performed to quantify the ability of Scouting HPS to find genuine τ_h candidates with p_{T} above 5 GeV. Figure 2 shows the results of this study. The efficiency of the Scouting HPS algorithm to reconstruct τ_h objects is shown as a function of the object's generator-level visible p_{T}. We note that the Scouting HPS algorithm is able to reconstruct τ_h objects at visible p_{T} down to 5 GeV. No previous τ_h reconstruction algorithm has shown any sensitivity in this region. Additionally, a marked falloff in the ability of the algorithm to identify τ_h objects at $p_{\mathrm{T}} > 30\,\mathrm{GeV}$, for all decay categories, is observed. This falloff is due to a tendency for high p_{T} charged hadrons to be reconstructed as e/γ candidates by the PF algorithm. Finally, we observe that most 1-prong, 2-strip τ_h that are found are reconstructed as 1-prong, 1-strip. This is a known phenomenon as the strip creation algorithm tends to merge π^0 decays. We note that for the first time we have demonstrated the ability to reconstruct τ_h objects at p_{T} below 20 GeV, with approximate efficiencies between 20% and 40% using the Scouting dataset.

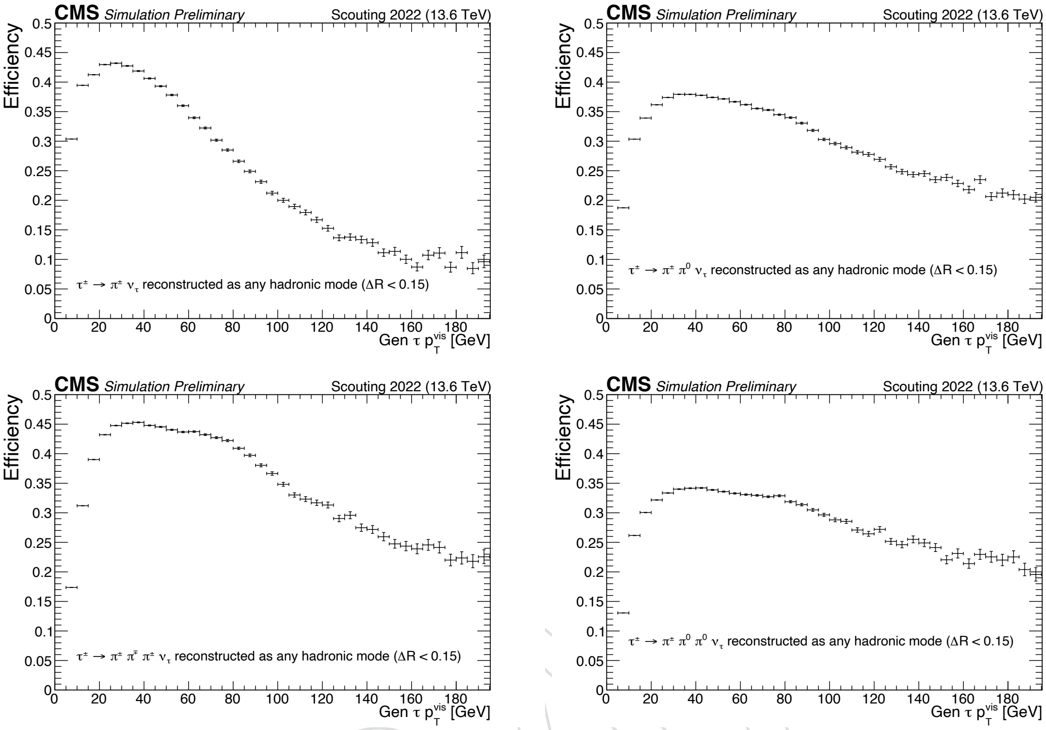

Figure 2: Simulation-level reconstruction efficiencies by τ_h decay category. A τ_h is considered to be found if Scouting HPS locates a candidate within $\Delta R < 0.15$ of the generator-level τ_h. The decay category of the ΔR matched candidate need not match the generator-level category.

2.2 TauNet neural network

The changes made in Scouting HPS drastically increase jet $\rightarrow \tau_h$ fakes. Traditional efforts to increase τ_h purity include isolation, multi-variate analysis (MVA) discriminants [8], and the DeepTau convolutional neural network [12]. Isolation only takes into account the p_{T} pollution within the isolation cone of the τ_h candidate, relying on the HPS algorithm to encode the structural information about the decay. MVA discriminants take into account more high-level information including lifetime, photon multiplicity, and ΔR information about strip constituents. The DeepTau neural network [12] uses a combination of high-level inputs and images of the decay constituents in the $\eta \times \phi$ plane available in the standard data stream. The Scouting τ_h reconstruction employs an Energy Flow neural network [13] to produce a likelihood discriminant using inputs described in detail in Section 2.2.1. Challenges present in the representation of composite particle objects such as the τ_h are that its constituents have no natural order and that there are a variable number of them. We sought a model that could accept the simple $\tau^{\pm} \rightarrow \pi^{\pm}\nu_{\tau}$ and complex $\tau^{\pm} \rightarrow \pi^{\pm}\pi^{0}\pi^{0}\nu_{\tau}$ equally well. An elegant solution to this is the Energy Flow architecture [13], which is based on DeepSets [14]. The defining feature of this architecture is its ability to map an unstructured set of inputs to a latent representation via summation. This representation is subsequently fed into a set of densely connected layers to finally calculate the classification likelihood. In our case, we consider the unstructured set to be the collection of Scouting PF particles within the AK4 jet containing the Scouting HPS τ_h. All constituents of the AK4 jet that seeds the τ_h are given to the model, as opposed to just those selected by Scouting HPS. This model is referred to as TauNet.

2.2.1 Network Inputs

TauNet takes in a 40×11 matrix where each column represents information about an individual PF candidate in the AK4 jet that contains the τ_h and each row represents an input feature. The strip constituents are "unpacked" from the strips. The input features of the particles include the $\Delta\eta, \Delta\phi$, and ΔR between the jet constituents and the jet axis. The hadron and e/γ entries may have values of either 0 or 1 depending on the species of the particle. The charge entry may only have values -1, 0, or 1. e/γ candidates are labeled as having charge 0. The d_{xy} (transverse impact parameter) and d_z entries only have non-zero values if the PF candidate had a track: charged pions and muons only. If a PF particle is identified as being a τ_h constituent by Scouting HPS, it is labeled 1 in its last column entry. Otherwise, spectator particles are labeled as 0. The number of allowed particles in each matrix is 40. If an input jet contained 15 particles, then the input matrix would contain 15 non-zero columns and 25 columns of zeros. The column order is shuffled randomly. Each non-discrete feature is transformed so that its mean is 0 and standard deviation is 1. Discrete features and all zero entries are left undisturbed.

2.2.2 Network Training

The training data are selected in the following manner. True Scouting τ_h candidates are selected if they are within $\Delta R < 0.15$ of their visible generator-level counterpart. Jets faking τ_h candidates are selected from same-sign $\tau_\mu\tau_h$ events in data. To reconstruct the τ_μ we use PF muon objects. We consider events with exactly one PF muon having $p_T > 3$ GeV. Great care is taken to make the inference of TauNet scale invariant. An $\eta \times p_T$ plane is divided into four tranches in η and 16 tranches in p_T. The η bins are spaced linearly with edges: [-2.4, -1.2, 0, 1.2, 2.4]. The p_T bin edges are logarithmically spaced from 5 to 60 GeV with no p_T limit on the last bin. In all, there are 4.6 million genuine τ_h candidates and an equal number of fakes in the training dataset. The receiver operating characteristic (ROC) distribution has an area under the ROC curve of 0.9.

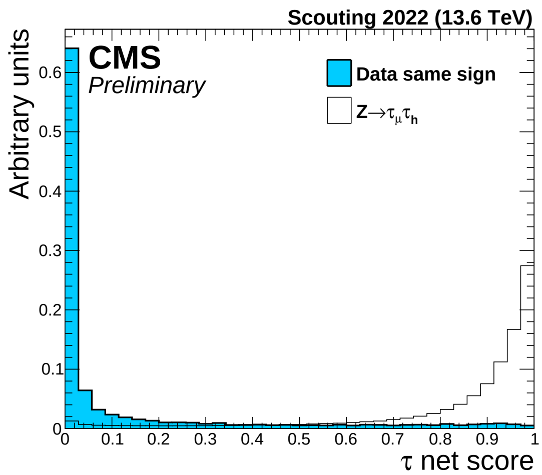

Figure 3: Distribution of TauNet score for a mass region containing $Z \rightarrow \tau_\mu\tau_h$ simulation. The filled histogram shows the TauNet distribution for same-sign data. Selections applied are outlined in the list below.

To illustrate the discrimination of TauNet for signal versus background, we show the distribution of TauNet for background and the simulated standard model Z boson in Figure 3. To maximize sensitivity to ditau resonances, the selection criteria below are applied after the TauNet score is determined.

- Selection criteria:
  - τ_h $p_T > 20\,\text{GeV}$
  - Muon $p_T > 20\,\text{GeV}$
  - μ isolation score > 0.9
  - Transverse mass $< 100\,\text{GeV}$ [8]
  - $(\text{ZetaMet} - 0.85 \times \text{ZetaVis}) > -25\,\text{GeV}$ [8]

2.3 Observation of $Z \rightarrow \tau\tau$ decays

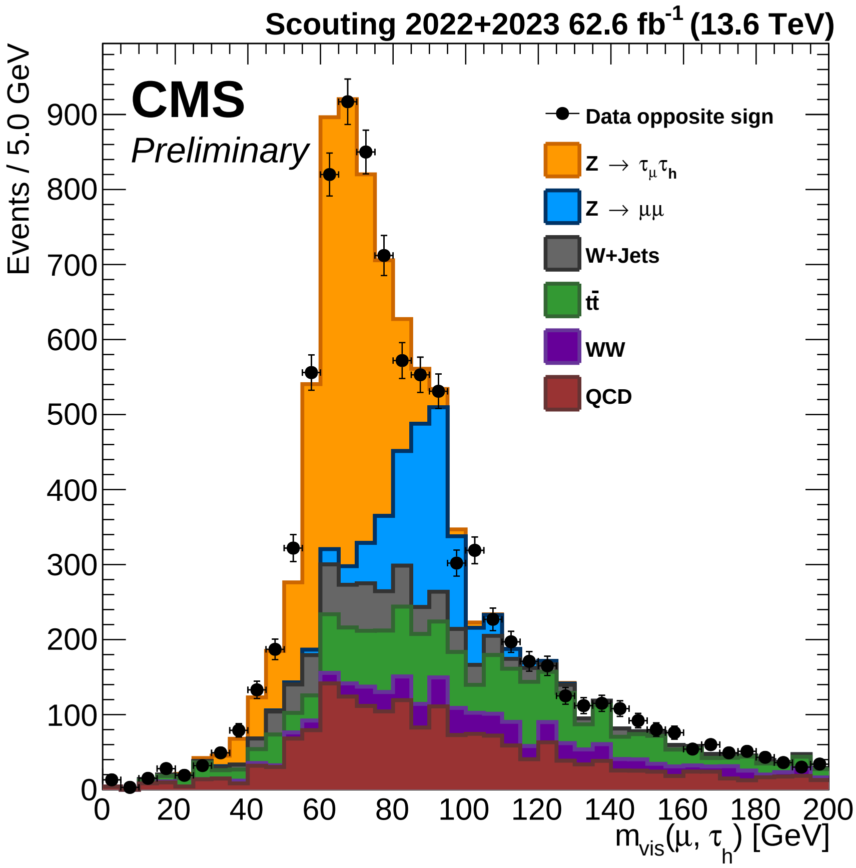

Figure 4: $Z \rightarrow \tau_\mu\tau_h$ visible invariant mass (m_{vis}) boson peak shown using the Scouting τ_h object on 2022 and 2023 data. The L1 trigger seeds used to collect this data required either a 30 GeV e/γ object, a single jet with p_T greater than 180 GeV, or an H_T greater than 360 GeV. Quality cuts are made to ensure that objects/quantities firing the L1 triggers are present. The extracted $Z \rightarrow \tau\tau$ cross section is 1835 ± 63 pb, within one sigma of the expected cross section of 1886 pb [7]. The QCD background is estimated by subtracting the same-sign components of all simulated backgrounds from the same-sign data component.

The 2022 and 2023 Scouting L1 seed menus did not include τ_h or single μ triggers [15]. Therefore, we proceed with a logical OR of L1 triggers that require a single isolated electromagnetic object above 30 GeV within $|\eta| < 2.1$, or $H_{\mathrm{T}} > 360$ GeV, or a single jet object with $p_{\mathrm{T}} > 180$ GeV. These triggers are described in [6]. To show the efficacy of the Scouting HPS and TauNet in combination to produce the Scouting τ, we present the Z boson peak in the $\tau_\mu \tau_h$ final state with combined 2022 and 2023 Scouting data. Selections applied to this data are listed in Section 2.2.2 with an additional requirement of TauNet score > 0.9.

Figure 4 shows the visible invariant mass of the τ_μ and τ_h. The filled circles show the data for which the τ_μ and τ_h objects have opposite sign. Backgrounds from $Z \to \mu\mu$, $t\bar{t}$, WW and W+jets are estimated from Monte Carlo simulation. The QCD background is estimated using data for which the τ_μ and τ_h objects had the same sign, from which the same-sign component of $t\bar{t}$, WW and W+jets have been subtracted to avoid double counting. The $Z \to \tau_\mu \tau_h$ distribution is also shown and used to determine the rate of this process. We use the branching ratio of τ decays to muons and hadrons to determine the number of $Z \to \tau\tau$ decays. The extracted $Z \to \tau\tau$ cross section is 1835 ± 63 pb, within one sigma of the expected cross section of 1886 pb [7].

3 Summary

The successful reconstruction of low-momentum τ_h decays using the Scouting dataset in Run 3 represents a significant advancement for the CMS experiment. Using an expanding cone algorithm allows the p_{T} threshold to be significantly lowered from the values in standard analyses. The TauNet model based on DeepSets achieves reasonable efficiencies and background rejection. Incorporating τ_h into the set of objects available in the Scouting dataset greatly extends the accessible phase space for BSM physics searches, particularly at low masses and transverse momenta. The meticulous development and validation of the Scouting τ_h reconstruction algorithm ensures accurate identification and measurement of τ_h candidates, as evidenced by the observed $Z \to \tau_\mu \tau_h$ process. This integration complements the existing suite of physics objects in the Scouting dataset. Consequently, it bolsters the overall discovery potential of the CMS collaboration in the quest for new physics beyond the standard model at Run 3.

References

[1] D. Curtin et al., "Exotic decays of the 125 GeV Higgs boson", *Phys. Rev. D* **90** (2014) 075004, doi:10.1103/PhysRevD.90.075004, arXiv:1312.4992.

[2] CMS Collaboration, "Searches for additional Higgs bosons and for vector leptoquarks in $\tau\tau$ final states in proton-proton collisions at $\sqrt{s}$ = 13 TeV", *JHEP* **07** (2023) 073, doi:10.1007/JHEP07(2023)073, arXiv:2208.02717.

[3] CMS Collaboration, "The CMS Experiment at the CERN LHC", *JINST* **3** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004.

[4] CMS Collaboration, "Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} = 13$ TeV", *JINST* **15** (2020) P10017, doi:10.1088/1748-0221/15/10/P10017, arXiv:2006.10165.

[5] CMS Collaboration, "Particle-flow reconstruction and global event description with the CMS detector", *JINST* **12** (2017) P10003, doi:10.1088/1748-0221/12/10/P10003, arXiv:1706.04965.

[6] CMS Collaboration, "Enriching the physics program of the CMS experiment via data scouting and data parking", 2024. arXiv:2403.16134. Submitted to *Physics Reports*.

[7] Particle Data Group, R. L. Workman et al., "Review of Particle Physics", *PTEP* **2022** (2022) 083C01, doi:10.1093/ptep/ptac097.

[8] CMS Collaboration, "Performance of reconstruction and identification of τ leptons decaying to hadrons and ν_τ in pp collisions at $\sqrt{s} = 13$ TeV", *JINST* **13** (2018) P10005, doi:10.1088/1748-0221/13/10/P10005, arXiv:1809.02816.

[9] M. Cacciari, G. P. Salam, and G. Soyez, "FastJet user manual", *Eur. Phys. J. C* **72** (2012) 1896, doi:10.1140/epjc/s10052-012-1896-2, arXiv:1111.6097.

[10] M. Cacciari and G. P. Salam, "Dispelling the N^3 myth for the k_T jet-finder", *Phys. Lett. B* **641** (2006) 57, doi:10.1016/j.physletb.2006.08.037, arXiv:hep-ph/0512210.

[11] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-k_T jet clustering algorithm", *JHEP* **04** (2008) 063, doi:10.1088/1126-6708/2008/04/063, arXiv:0802.1189.

[12] CMS Collaboration, "Identification of hadronic tau lepton decays using a deep neural network", *JINST* **17** (2022) P07023, doi:10.1088/1748-0221/17/07/P07023, arXiv:2201.08458.

[13] P. T. Komiske, E. M. Metodiev, and J. Thaler, "Energy Flow Networks: Deep Sets for Particle Jets", *JHEP* **01** (2019) 121, doi:10.1007/JHEP01(2019)121, arXiv:1810.05165.

[14] M. Zaheer et al., "Deep sets", *Advances in neural information processing systems* **30** (2017) arXiv:1703.06114.

[15] CMS Collaboration, "Development of the CMS detector for the CERN LHC Run 3", *JINST* **19** (2024) P05064, doi:10.1088/1748-0221/19/05/P05064, arXiv:2309.05466.